Legal knowledge based systems

JURIX 92

Information Technology and Law

The Foundation for Legal Knowledge Systems

Editors:

C.A.F.M. Grütters

J.A.P.J. Breuker

H.J. Van den Herik

A.H.J. Schmidt

C.N.J. De Vey Mestdagh

J.A.P.J. Breuker, On Legal Information Serving, in: C.A.F.M. Grütters, J.A.P.J. Breuker, H.J. Van den Herik, A.H.J. Schmidt, C.N.J. De Vey Mestdagh (eds.), Legal knowledge based systems JURIX 92: Information Technology and Law , The Foundation for Legal Knowledge Systems, Lelystad: Koninklijke Vermande, pp. 93-102, 1994 ISBN 90 5458 031 3.

More information about the JURIX foundation and its activities can be obtained by contacting the JURIX secretariat:

Mr. C.N.J. de Vey Mestdagh
University of Groningen, Faculty of Law
Oude Kijk in 't Jatstraat 26
P.O. Box 716
9700 AS Groningen
Tel: +31 50 3635790/5433
Fax: +31 50 3635603
Email: sesam@rechten.rug.nl

 http://jurix.bsk.utwente.nl/

ON LEGAL INFORMATION SERVING

J.A.P.J. BREUKER
Department of Computer Science and Law, University of Amsterdam, Amsterdam, The Netherlands

Summary

"Can I park on the left hand side?", or "Must I declare stolen property?" are two legal questions that are intriguing by their phrasing and their contents. Although it is suggested that modern (intelligent) database technology would be an important answer to the increasing need for legal information and advice, it is still an open question to what extent this is applicable. In this paper I argue that the applicability of the technology is very limited, because it presupposes that such a question has some legal facts (a set of applicable rules, or cases) as its answer. It has not. Answering legal questions -also called legal information serving- involves making inferences about the world and about the structure of a regulation, as can be shown by analysing these questions. Legal information serving is more complex than typical problem solving, because no fixed, domain-independent methods appear to be used. However, in a legal reasoning architecture where reasoning about the world is not mixed with reasoning about applying the regulations, a non-combinatorial, tractable solution can be found. This is the more so, when reasoning about the world can be limited to its superficial, terminological structures.

1 . Introduction

In his survey on legal information processing systems, Sergot [1991] distinguishes two major types:

1. Storage and retrieval of legal sources
2. Legal analysis systems

The first category is traditionally supported by database technology. Aside from the general office automation tools (text processing; databases for office procedures), databases for legal sources are probably the most widely used tools in legal practice. In the Netherlands, Vermande's NLEX or PARAC are examples containing legislation. Kluwer's Legal Database contains precedence cases. These are traditional databases, allowing access via keywords or full text. I will not digress on the problems that are inherent in this way of information handling [e.g., Mital et al., 1991]. They can be summarised by the fact that even intelligent tools for navigating through such (large) data bases will easily lead to getting lost. One can get lost by retrieving too much information as a consequence of the use of unspecific terms. More problematic is the case where no information is returned. Negation by failure, i.e., by the fact that nothing is found, is not possible. The required information may be there, but the queries may have been inappropriate. These problems are typical for all types of (large) databases. However, for retrieving information on legal sources, it appears that database technology is largely in insufficient, if not inappropriate at all, as I will argue in section 2. It turns out that typical requests for legal information require inference capacity. Moreover, controlling the use of this inference capacity appears to be more complex and varied than in problem solving (section 3), and needs dialogue support (section 4). In section 5 a framework for legal information serving is presented.

2 . Requesting legal information

As the following examples of legal requests may show, often there is not a direct textual or conceptual correspondence between the topic and the legal source.

1. Does the law say anything about declaring stolen property?
2. Should I marry to avoid being expelled?
3. Is it forbidden to borrow a house?
4. Can I park my car on the left hand side of the road?

Aside from the fact that specific relationships are requested, a data base retrieval will in all cases return many sources, because the terms used are suffiently general. However, not one of these sources will contain an answer to these questions. Moreover under the naive rule that everything is allowed that is not forbidden, or prescribed in a specific way, the user is tempted to conclude, that he will not get into conflict with the law if no matching legal source is found. These questions may look typical layman's requests for advice, but they are not different from questions a professional lawyer may have in mind, when exploring domains of law with which he is not familiar.

Let us examine in more detail what the last question (4) entails. In the preambule to the new Dutch traffic code (RVV-90) it is suggested that in contrast with to the old code (RVV-66), it is now permitted to park on the left hand side of the road, because current the code contains only statements where parking is forbidden. No explicit statement on left hand side parking is to be found in the code itself. A search using the term "left hand side" will yield nothing; the term "parking" will produce at least five articles. However, no answer is to be found in these sources. The preambule holds probably the same view as the layman by assuming that if it is not forbidden, it is allowed. In practice it is not permitted to park on the left hand side of the road because in trying to do so, one will violate a rule that seems to have nothing to do with parking. In particular, the general rule that vehicles should keep the right hand side of the road (art. 3.1.) appears to prevent any parking on the left hand side, although it does not contain these terms explicitely. Keeping to the right implies that a parking action to the left is excluded, or that "to keep" also subsumes "to park".

Direct search fails because regulations contain terms which imply other terms. These implications are not always straightforward subsumption relations. The relations may be subtle and complex. For instance, to see that article 3.2: "A pair of bicycles is permitted to ride next to one another" is an exception to 3.1., one has not only to establish the subsumption of bicycles as vehicles, but also that in the spatial arrangement indicated by "next", the left hand bicycle of the pair is not keeping to the right. In other words, spatial relations have to be derived in order to understand the relationship between the two articles.

It appears that the structure of a regulation is not explicit in the regulation itself: only by spatial reasoning one can establish the fact that art. 3.2. is an exception to art. 3.1. Note that it is not a full exception, because the left hand bicycle is not permitted to ride on the middle of the road: it should still keep as right as possible. Only a very global, and inaccurate structure is explicit in the code itself. Regulations are drafted in such a way that the more general rules or statements are followed by[1] exceptions or further specifications. In matching the request with a regulation, all articles have to be checked whether they contain exceptions, because these exceptions may be relevant for the request in an *indirect way*. For instance, there are a number of exceptions to the keeping to the right rule (art. 3.1.); article 3.2. being one of these. However as the question (4) is about a car, and it can be inferred that a car is not a bicycle -or that riding is not parking- one can decide that this exception is not relevant for our question. Checking these potential exceptions, it seems that left hand parking is not allowed: not even in one way streets.

Whether an article contains a relevant exception or modification is quite often not aparent without making inferences about the world. For instance, the equivalent question posed in Great Brittan - "Can I park my car on the right hand side of the road" -, should be answered with: "Yes, (under a list of conditions), but only during the day, or in a one way street". This is due to the fact that the British traffic code contains an article which states that the reflector of a car should be visible at night for the (left hand) driving traffic.

The upshot of this short and informal analysis of "what's in a legal question?" is that answering almost always requires making inferences. They are of two types: about the world the regulations refer to, and about the relationships between articles (or cases) in a legal source, in particular the structure of exceptions. This full answering of legal questions is called "*legal information serving*", in contrast with database retrieval aimed at collecting legal sources which contain similar or the same terms as the request.[2]

3. Information serving and problem solving

3.1. Problem solving methods for controlling inferences

As the examples also suggest, legal information serving has both the flavour of question answering and of problem solving. "Can I park my car on the left hand side of the road?" looks like a simple, factual question, which should yield a simple yes/no answer with some potential elaboration: e.g. "Yes, if there is a parking strip". However, the question: "what kind of contract should I propose to prevent my students from paying copyright to my publisher" looks rather as asking for a legal advice. It involves legal problem solving, appearing as a typical design problem, in which the question serves as a design constraint.

The distinction between legal information serving and legal problem solving is gradual. It is not dependent on the nature of the request. As we have seen, the same type of request may lead to completely different ways of finding the answer. It is dependent on the structure of the legal knowledge, and on the availability of *problem solving methods.* These two factors are not independent, as the "interaction hypothesis" suggests [Chandrasekaran, 1987]. The interaction hypothesis states that the structure of knowledge is geared to, or allows specific problem solving methods. A problem solving method is either a decomposition of a problem solving task, or a procedure (algorithm) to solve a problem [Steels, 1990]. A problem solving method therefore *controls* the use of inferences by stating (alternative) sub-goals. These decompositions are not arbitrary. In a good decompostion, the process of inference making becomes tractable and does not explode, because only a subset of all potential inferences is to be executed. The subgoals act as short range, intermediary aims, where the final aim is not visible from the starting point. It should be noted that the availability of matching problem solving methods is what gives human and much artificial problem solving a heuristic flavour, and not the nature of the domain knowledge.

3.2. No general methods for solving legal requests.

However, in legal information serving most often no specific, apropriate problem solving methods are available. For instance, there is no way to decompose the question on parking in such a way that restricted parts of the knowledge, i.e., the RVV-90, can be used. In solving this problem one has to search the code several times almost exhaustively, and a global insight into its structure is not sufficient to plan an efficient strategy. Without a problem solving method, the inferences become undirected and intractable. The structure of general and specific (exception) articles can lead in the first place to fastly branching search. However, the stopping problem makes it even worse. The search can only halt by a match or a contradiction, and both are often not derivable

from the regulation itself, but by common sense reasoning (e.g., the required spatial reasoning on the pair of bicycles).

To give some idea how fast things may explode the parking example may give some indication. The first thing to look at is the article on parking (RVV-90, art. 26). This is a list of five conditions under which parking is not allowed, e.g., at a crossing. For each of these conditions one can check whether it entails "left hand parking", i.e., whether there is a left hand side. A crossing is a road and therefore has a left hand side. In four of the five conditions there is a "left hand side of a road" implied, which would yield as an answer: "left hand side parking is at least not possible at crossings, along a yellow, continous line, etc." Checking all five conditions yield in the long run indeed contradictions, so search along these lines can be stopped (I will leave out all details here). However, the famous article 3.1. which states that vehicles should keep to the right appears to forbid any left hand parking, assuming that "keeping" subsumes "parking", i.e., there is a match. The next step is to check whether there are exceptions to this article, and indeed there are at least ten -for the human interpreter. For instance, on a roundabout we are permitted to keep the left lane, and assuming that a roundabout is a crossing of roads, it does not allow for parking (contradiction). Without a specific representation and search method, a precise number of steps cannot be given. The TRACS system, that is capable of making these inferences is still incomplete [Haan, 1992].

3.3. Falsification by exclusion: the use of taxonomies

Most of these chains of inference are stopped by contradictions on the basis of exclusive types: e.g., a car is not a bicycle. The use of explicit taxonomies in definitorial law (e.g., art. 1. of the RVV-90) is therefore of extreme importance, both for human and artificial legal reasoners. For instance, it is not clear whether "keeping to the right" also entails "parking to the right". If it does, no left hand parking appears to be possible; also in a one way street. If it does not, then parking on any part of the road is possible. This makes the application of many sorted logics in artificial legal reasoning very profitable, because these logics contain built-in algorithms for falsification by exclusion. On the other hand there is still the knowledge acquisition problem to (re)construct these taxonomies. Legal sources are notably sloppy or implicit on this information.

To avoid misunderstandings, this has nothing to do with the fact that regulations may contain open concepts to leave room for bottom-up interpretation (precedence). Also vague and abstract terms should be intensionally explicit, i.e., their relation to other terms used in a regulation should be stated if a relation is assumed. The vagueness or openess of terms is associated with the fact that the *referents* in the real world may be difficult to identify. Taxonomies are a minimal condition for the coherence that a legal source should exhibit [Brouwer, 1990]. Taxonomies are not the only source for exclusion. For instance, in spatial reasoning two (solid) objects cannot be in the same position (cf the pair of bicycles rule).

In summary, the inferences in legal reasoning are combinatorial, and can be very long, because the stopping conditions are not always (explicitely) available. The use of problem solving methods (task decomposition) can keep both problems within limits by providing subgoals, i.e., states which function both as pragmatic stopping conditions, and heuristics for selecting plausible alternatives in reasoning.

3.4. In search of legal reasoning strategies.

I have somewhat exagerated when stating that there are no task decompositions in legal information serving. What I meant was that they are not of the same type as in typical problem solving tasks, as diagnosis, monitoring, planning, design, etc., where relatively domain-independent, stable methods can be found. In legal reasoning these decompositions appear to be (also) dependent on the structure of the legal source. Legal experts who have a good insight into the structure of dependencies in a legal source, can easily separate out potentially irrellevant articles. They may start with translating the terms in the request into the appropriate legally defined ones, i.e., identifying these in the taxonomies. Currently we (Muntjewerff, Den Haan) are investigating what other sources for decompositions may occur by studying thinking aloud protocols of practicing experts in the Narcotics Act (Opiumwet). Thusfar it is clear that experts make their decompositions on the basis of the structure of the source, while legal practitioners (lawyers), who are not specialists, try to apply general schemas (phases of establishing "the facts", finding relevant sources, etc.) by alternatively exploring the sources, and reflecting on next courses of action based upon general characteristics of the legal source. In other words, the decomposition is constructed dynamically by strategic reasoning. [Breuker & Wielinga, 1989]

4. Intention and cooperation in legal information serving

Problem solving strategies can cut down the search space in legal information serving. However, this search space is in the first place dependent on the nature of the request. The more specific and unambigious the request, the smaller this space. In human and also in artificial advice this is accomplished by the use of dialogue techniques.

The example questions in section 2 may look somewhat exotic. Who would ask such questions, and particularly why? By providing a *context,* the meaning of the questions becomes more obvious and specific. For instance, the question whether a house can be borrowed (3) should be taken as in contrast with buying or renting one. The context may be that the client wants to evade the House Renting Act (Huurwet), because it contains articles that make terminating a contract difficult for the owner. Therefore she may want to know whether lending is permitted, and whether this lending indeed evades this Act, because a lender(-owner)/borrower relationship has not been foreseen by it. The car parking question (4) can be part of a case in which the client is in a legal dispute about a fine or an accident. In that case the question can be more specific. For instance, if it turns out that the car was parked outside city limits on a main road (art. 25c, RVV-90), the question is easily answered. In general, when the request is about a specific case, information serving is not much different from assessing a legal case, for which problem solving methods have been identified [Schreiber et al., 1993]. In legal advice, the context and specification of the request -intentions- of the client is obtained in a dialogue. The legal practitioner tries to obtain either a more specific goal, and/or a number of intermediary goals by "negotiating" with the client. In that respect the dialogue can yield a somewhat similar function as the problem solving methods.

Recognizing the intentions of users/clients is a notably hard and interesting problem in modelling dialogues. Research on problems in database retrieval has been the major source of inspiration to clarify the problem and to propose solutions in the form of user modelling [Kok, 1990], and dialogue mechanisms which allow detailed interpretation of the user's request. [McKeown, 1985][Pollack, 1986][Cohen et al.,1989] These intelligent dialogue interfaces to databases can be viewed as advanced types of "conceptual" legal database retrieval systems. [Bing, 1987]

5. Towards a legal information server

What appeared as a simple function - retrieval of relevant legal sources - turns out to be very complex and almost impossible to automate, or to support by an artificial legal reasoning system. Not only database technology appears inadequate, the analysis also shows that solutions are complex and probably beyond the state of the art of current AI research. I should have added: principled solutions, because in AI it is often the case that the difference between a tractable and an intractable solution is very small. Can we design such solutions for legal information serving without sacrificing too much functionality, or incomplete or erroneous solutions?

5.1. Reasoning about the world and about regulations

The solutions proposed are part of the architecture for legal reasoning as described earlier [Breuker & Haan, 1990][Breuker, 1991]. Central to it is that regulation knowledge, i.e., the knowledge contained in a statutory legal source, is separated from the knowledge about the world to which these regulations refer. This leads to a decomposition of legal reasoning in two different components: one that is concerned with reasoning about the world, and the other one about legal consequences, i.e., the normative aspects of the law.

The first component makes inferences about the world to which the regulation refers. It is a model of how this world behaves and represents the views the legal system has about some subsystem in society [Valente & Breuker, 1991][also Breuker & Valente, 1992]. This is expressed by definitorial statements in legal sources. The taxonomies discussed above are an essential part of the description of this world. However, often this knowledge is implicit and has to be derived from the terms used in the regulations and from supporting legal documentation (e.g., memories). For each legal domain such a world has to be (re)constructed, and represented. This is called the WORLD MODEL and supports the understanding of states and events in the domain. For instance deriving the fact that "next" in "a pair of bicycles next to one another" means that one is less to the right side of the road than the other one, is performed here. The WORLD MODEL is used to construct a SITUATION DESCRIPTION, which is a situation specific model of the case or request at hand. [Steels, 1990] It consists of a list of states and events in terms that are known to the legal reasoning component.

This is the second component, and its role is to derive the legal consequences of the SITUATION DESCRIPTION. These legal consequences can be in terms of violations of regulations, or prescriptions for actions according to the law. This component consists of two sub-components. The first one applies the regulations as the are represented in the REGULATION KNOWLEDGE BASE. The second sub-component resolves the conflicts that occur when applicable regulations appear to contradict one another in their consequences. This occurs, e.g., when two rules are both applicable to the same subset of states in the SITUATION DESCRIPTION, and have different consequences. One rule is then the exception to the other one, and these conflicts are resolved by using legal meta-knowledge (e.g., the "lex specialis" rule). The legal reasoning component can be realised as a classical production system with knowledge based conflict resolution, and enhanced with taxonomic unification.

The decomposition of legal reasoning into reasoning about a legal world and about legal consequences, is a built-in task decomposition which reduces the complexity considerably. The legal reasoning component is utterly tractable, because the unification of condition parts of the rules and the description of the situation is a taxonomic one, as in sorted logic. If a rule is about "vehicles" and the situation description contains the term "car", the rule matches. In general legal rules have as their consequences a statement that a state or action is obliged or forbidden, but other "modalities" may occur also (e.g., definitions, prescriptions). These complications have no combinatorial effects, and only affect processing time in a linear way.

This decomposition makes the application of regulations completely tractable. Taxonomic (sorted) unification combined with straightforward rule application and conflict resolution can handle legal reasoning in an almost non-combinatoric way; the only combinatorics involved are related to combinations of terms in rules. As an emerging property, (hidden) dependencies between rules can be traced, i.e., which rules are exceptions or modifications to which other rules, and in which circumstances.

This method, or task decomposition, appears to be domain independent. It works for the traffic domain [Haan, 1992}, and it is being tried out in other legal domains (narcotics; intellectual property). It is certainly not the strategy employed by human legal experts or novices. Most likely the reason is that it would lead to overflow of short term memory in humans. Keeping track of a list of applicable and contradictary outcomes is very difficult. Humans prefer to resolve all these conflicts immediately to keep a focus on the problem solving topics, because that keeps the short term memory load as low as possible. The computer has no real limits in maintaining a short term memory, but is far less efficient in recognizing patterns. In this respect the proposed architecture is really an *artificial* legal reasoner. A second problem for human problem solvers is that they know far more about the world than what the law supposes. It is hard to distinguish between common sense reasoning and trying to conceive the world as the legislator has viewed it, because it is rather abstract with respect to the specific cases. This distance between abstraction and real case descriptions, combined with the fact that these abstractions are largely implicit, makes legal reasoning very complex. These world models are models of subsystems in society. Both society and the views on how these subsystems should be controlled change. Therefore, the correspondence between the world model and the regulations may become less direct. In general, this leads to modifying the applicability of regulations by precedence when the regulations themselves are not easily, or transparently updated. Aside from the fact that regulations may also be the result of political compromises which leave ambiguous interpretations of the world, it is obvious that the major problem in constructing an artificial reasoner for some legal domain is not the representation of the regulations, but of the world it refers to. The RVV-90 is a relatively simple domain, because it (also) includes current practice and debate on traffic safety.

5.2. Abstracting from reasoning about the world

What has remained implicit fact that probably not a full world model is required, but only the components from which it is built: the concepts or terms. In a legal information serving the interest is not in knowing how a particular world (subsystem) works, i.e., what behaviours can be derived from this model, but only to interpret behaviours (states, events) which are specified by the user.

Therefore, the focus appears to be rather on the classificatory aspects of what the world contains, than on the fact that one can "animate" (simulate) this world. This taxonomic view on the model of a domain coincides with theories about knowledge representation, where *terminological* (taxonomic) formalisms like the KL-ONE family are distinguished from "axiomatic" representations (e.g., logics). The terminological view represents the ontology of a domain, where the axiomatic view provides the corresponding calculus to make the inferences about statements (facts) in the appropiate terms. The question is whether we need all these axiomatic inferences about the world in legal reasoning.

In the current architecture, taxonomic representations act as the interface between the legal application component and the world component. However, once a situation is specified, there is no longer the need to derive behaviour from the world model, and only taxonomic information is used to decide on the applicability of rules. The full world model is needed to derive values or states which were not part of the specification of the input. However, that is not the case when the user comes with a request that is already specified in terms of a situation description of a legal domain. In other words, when we present the user with in interface which allows him to specify his request only in the

taxonomic terms (and structure) which the legal domain presupposes, we can avoid also the reasoning about the world. This kind of interface is not difficult to conceive. The user can select from menus or palettes the terms for composing a request. These terms are directly "understood" by the legal application component. For instance, in the TRACS system which assesses the legal consequences of the RVV-90, the user interface allows the user to specify in a graphical way a situation [Haan, 1992]. This graphical situation maps directly on the underlying taxonomy of the RVV-90, so that reasoning about the world of traffic is not really required; only about their taxonomic relations. The reason one can minimize the world (common sense) reasoning to terminological views, is that in information serving one does not have to assess, and decide a case. For instance, one does not have to establish the fact whether an agent has, e.g., signed a contract without being forced to do so. It is assumed that the user will provide this information.

The observation that terminological knowledge about a legal domain, plays a pivotal, and probably exclusive role in legal information serving, makes the amount of inferencing rather tractable. However, the fact that, e.g., for the traffic regulations (RVV-90) *spatial* reasoning is required, seems at odds with this observation. Indeed full spatial reasoning implies a calculus, which can derive positions and spatial relations of objects. However, in TRACS we do not perform full spatial reasoning. Spatial knowledge is "compiled out" into relation terms - e.g., next, left_of, right_of, etc. - for which the meaning (implications) are obtained by procedural attachment to a terminological entry. For instance, in the taxonomic representation "left_of" excludes "right_of". However in the procedural attachment, the inverse relationship between the two can be established. Procedural attachment in terminological representations are not very principled, but very practical.

6. Conclusions

The major thesis of this paper is that legal information serving cannot be supported by database technology, because it is not fact retrieval, but deriving legal consequences from situations. At first glance, it looks more complex than typical problem solving tasks, because it lacks domain independent methods for controling the inferences. However, a decomposition in terms of reasoning about the legal world and of deriving the legal consequences of a situation, makes the legal reasoning probably tractable. At least the derivation of legal consequences can be performed by sorted unification, and rule application with conflict resolution to trace exceptions.

The next "trick" proposed is to leave the reasoning about the world to the user by providing a user interface that represents the taxonomic views on the model of the world. This simplification may not work completely, because types of reasoning may be applicable that cannot be captured by a terminological framework. In that case, this knowledge - e.g., about spatial relationships - may be compiled out and associated with the terms by procedural attachment.

If taxonomic knowledge is an interface between world knowledge and regulation knowledge, it appears that in legal information serving can start there. It also solves -or rather: avoids- problems of communication by the use of dialogue with the user because the user is restricted now to formulate his requests in terms of the law. This may look like a straightjacket, but in most cases it will be viewed as a help to formulate a request. However, for this simple solution, there is also a price to pay. The description of the situation may consist of combinations of terms that do not make sense at all. For instance, nothing prevents the user from specifying that a tram parks on a pedestrian crossing. A full interpretation by applying world knowledge would prevent such physical anomalies.

Indeed, by restricting the inference and communication capabilities to obtain pragmatic solutions I seem to be back at the notion of a "conceptual front end" [Bing, 1987] for

legal information serving. However, this "conceptual front end" that is driven by taxonomical knowledge, is not built to communicate with a database, but with an enhanced production system, and a somewhat hybrid terminological representation formalism.

7. Notes

1 A design principle in drafting regulations is to start with the most general statements. Legal draftsmen may view this order of articles as corresponding to the "lex specialis" principle. However, it may only approximate, because a one dimensional sequence is not suited for expressing a potentially multidimensional structure of exceptions and amendments to the general rules [Haan, 1992].

2 Information serving appeared as a newspeak term for intelligent database technology in the latest call for proposals in the Esprit-III programme.

8. References

This article may only be cited as: Breuker, J.A.P.J., On Legal Information Serving. In: Grütters, C.A.F.M., J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh (eds), *Legal Knowledge Based Systems: Information Technology & Law, JURIX'92*, Koninklijke Vermande, Lelystad, NL, 1992.

[Bing, 1987] Bing, J., Designing text retrieval systems for "conceptual searching". In: *Proceedings of the First International Conference on Artificial Intelligence and Law*, Boston, 1987, pp. 43-51.

[Breuker & Haan, 1991] Breuker, J.A.P.J. and N. den Haan, Separating world and regulation knowledge: where is the logic? In: Sergot, M. (ed), *Proceedings of the third international conference on AI and Law*, New York, Association of Computing Machinery, 1991.

[Breuker & Wielinga, 1989] Breuker, J.A.P.J. and B. Wielinga, Model Driven Knowledge Acquisition. In: Guida, P. and G. Tasso (eds), *Topics in the Design of Expert Systems*, Amsterdam, North Holland, NL, 1989, pp. 265-296.

[Breuker, 1990] Breuker, J.A.P.J., Towards a Workbench for the Legal Practitioner. In: Noortwijk, C. van, A. Schmidt and R. Winkels (eds), *Legal Knowledge Based Systems: Aims for Research and Development, JURIX '90*, Koninklijke Vermande, Lelystad, 1990.

[Brouwer, 1990] Brouwer, P.W., *Samenhang in Recht*, Wolters-Noordhoff, Groningen, 1990.

[Chandrasekaran, 1987] Chandrasekaran, B., Towards a functional architecture for intelligence based on generic information processing tasks. In: *Proceedings of the 10th IJCAI*, Milano, IT, 1987, pp. 1183-1192.

[Cohen et al., 1990] Cohen, P.R., J. Morgan and M.E. Pollack, *Intentions in Communication*, MIT Press, Cambridge, MA, 1990.

[Haan, 1992] Haan, N. den, TRACS: a support tool for drafting and testing law. In: Grütters, C.A.F.M., J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh (eds), *Legal Knowledge Based Systems: Information Technology and Law, JURIX '92*, Koninklijke Vermande, Lelystad, NL, 1992.

[Kok, 1990] Kok, A.J., *User Modelling for Data Retrieval Applications*, PhD Thesis, Universiteit van Amsterdam, NL, 1990.

[McKeown, 1985] McKeown, K.R., Discourse strategies for generating natural language text, in: *Artificial Intelligence*, 1985, no. 27, pp. 1-41.

[Mital et al., 1991] Mital, V., A. Stylianou and L. Johnson, Conceptual information retrieval in littigation support systems. In: Sergot, M. (ed), *Proceedings of the third conference on AI & Law*, ACM, New York, 1991, pp. 235-243.

[Pollack, 1986] Pollack, M.E., *Inferring Domain Plans in Question Answering*, PhD thesis, University of Pennsylvania, 1986.

[Schreiber et al., 1993] Schreiber, G., B. Wielinga and J.A.P.J. Breuker (eds), *KADS: a principled approach to knowledge based system development*, Academic Press, New York, 1993.

[Sergot, 1991] Sergot, M.J., The representation of law in computer programs: A survey and comparison of past and current projects. In: Bench-Capon, T.J.M. (ed), *Knowledge based Systems and Legal Applications*, Academic Press, 1991.

[Steels, 1990] Steels, L., Components of Expertise, in: *AI Magazine*, Summer 1990, pp. 29-49.

[Valente & Breuker, 1991] Valente, A. and J.A.P.J. Breuker, Law Functions: modelling principles in legal reasoning. In: Breuker, J.A.P.J., R.V. de Mulder and J.C. Hage (eds), *Legal Knowledge Based Systems: Model Based Legal Reasoning, JURIX '91*, Koninklijke Vermande, Lelystad, NL, 1991, pp. 40-52.

[Valente & Breuker, 1992] Valente, A. and J.A.P.J. Breuker, A Model-Based Approach to Legal Knowledge Engineering. In: Grütters, C.A.F.M., J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh (eds), *Legal Knowledge Based Systems: Information Technology and Law, JURIX '92*, Koninklijke Vermande, Lelystad, NL, 1992.